DOCUMENT RESUME

ED 474 048 SE 066 927

AUTHOR Ryan, David M.

TITLE Mathematics in Action: Two New Zealand Case Studies.

PUB DATE 2000-00-00

NOTE 19p.; In: Proceedings of the International Conference on Technology in Mathematics Education (Auckland, NZ, December 11-14, 2000). For full proceedings, see SE 066 920.

AVAILABLE FROM Aukland University of Technology, PB 92006, Auckland, NZ. Tel: 64-9-907-999 X 8405; Fax: 64-9-307-9973; e-mail: mparker@aut.ac.nz or University of Auckland PB 92019, Auckland, NZ. Tel: 64-9-373-7599 X 5886; Fax: 64-9-373-7457; e-mail: minlee@math.auckland.ac.nz.

PUB TYPE Reports - Research (143)

EDRS PRICE EDRS Price MF01/PC01 Plus Postage.

DESCRIPTORS Case Studies; Foreign Countries; Higher Education; Industry; Innovation; *Mathematical Applications; *Mathematical Models; Mathematics Instruction; *Relevance (Education)

IDENTIFIERS Discrete Mathematics; *New Zealand

ABSTRACT

Mathematics is playing an increasingly important role in business and industry. In this paper we present two case studies to illustrate the power and impact of mathematics in two important practical applications in New Zealand. The first case study describes the development of a mathematical optimization model to maximize the value of aluminum produced at New Zealand Aluminum Smelters Ltd. The second case study describes the development and implementation of state of the art optimization models and solution methods to solve all aspects of the crew scheduling problems for Air New Zealand Ltd. (Author)

Reproductions supplied by EDRS are the best that can be made from the original document.

PERMISSION TO REPRODUCE AND DISSEMINATE THIS MATERIAL HAS BEEN GRANTED BY

M. Thomas

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

1

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☑ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

MATHEMATICS IN ACTION: TWO NEW ZEALAND CASE STUDIES

David M. Ryan
University of Auckland
<d.ryan@auckland.ac.nz>

Mathematics is playing an increasingly important role in business and industry. In this paper we present two case studies to illustrate the power and impact of mathematics in two important practical applications in New Zealand. The first case study describes the development of a mathematical optimisation model to maximise the value of aluminium produced at New Zealand Aluminium Smelters Ltd. The second case study describes the development and implementation of state of the art optimisation models and solution methods to solve all aspects of the crew scheduling problems for Air New Zealand Ltd.

ED 474 048

Introduction

During the past fifty years the power and relevance of mathematics in business and industry have grown in both the variety of applications and the importance of its impact. Besides the development of the underlying mathematical techniques, an obvious reason for this exciting trend has been the development of computer technology and the associated computational techniques. In particular methods of mathematical optimisation including linear and integer programming have become widely used to model and solve many important practical problems. Scheduling and resource allocation decision problems occur in many business and industrial organisations. Often these problems involve valuable or scarce resources such as time or materials or people – finding optimal or near optimal solutions of these problems can provide millions of dollars of savings and provide a significant competitive business advantage.

The general linear programming (LP) model with m constraints and n variables has the form

$$\begin{array}{lll} \text{LP:} & \text{minimise} & z = \mathbf{c}^T\mathbf{x} \\ & \text{subject to} & A\mathbf{x} = \mathbf{b} \\ & \text{and} & \mathbf{x} \geq \mathbf{0} \end{array}$$

where A is an $m \times n$ real matrix and $\mathbf{b}$ and $\mathbf{c}$ are given real vectors of dimension m and n respectively. The mathematics of linear programming and the associated solution methods are interesting in their own right and in fact students taking Mathematics with Statistics at secondary school study linear programmes in two variables. Even beyond the beauty of the mathematics of LP, the relevance of the LP model in many practical applications makes linear programming one of the most important mathematical developments in the 20th century.

As a very special case of the general LP model, the set partitioning model provides an underlying mathematical model for many scheduling applications. The set partitioning problem (SPP) is a specially structured zero-one integer linear programme with the form

$$\begin{array}{lll} \text{SPP:} & \text{minimise} & z = \mathbf{c}^T\mathbf{x} \\ & \text{subject to} & A\mathbf{x} = \mathbf{e} \\ & \text{and} & \mathbf{x} \in \{0,1\}^n \end{array}$$

SE019927

BEST COPY AVAILABLE

where $\mathbf{e} = (1,1,1,\ldots,1)^T$ and A is a matrix of zeros and ones and the solution vector $\mathbf{x}$ must take values of zero or one. Because of the computational difficulties in solving very large and practical instances of the set partitioning problem, many early attempts to use optimisation solution methods to solve scheduling problems were unsuccessful and researchers resorted to a variety of heuristic solution methods. While most heuristic methods are relatively easy to implement and may have reasonably inexpensive computer resource requirements, they suffer from two major disadvantages. Firstly they can provide no bound on the quality of any feasible solution that they produce, and secondly they are unable to guarantee a feasible solution will be found even if one exists. The heuristic methods may fail to find a feasible solution either because the heuristic method is inadequate or because the problem is truly infeasible. In contrast, an optimisation method can reliably detect infeasibility. During the past two decades, the development of optimisation methods and techniques for the solution of set partitioning problems, and the increase in computer power has meant that we are now able to solve realistic-sized models that arise in many practical scheduling problems.

In this paper, we present two case studies to illustrate the power and impact of mathematics in two important practical applications in New Zealand.

Case 1: Optimised Cell Batching at New Zealand Aluminium Smelters Ltd

New Zealand Aluminium Smelters Ltd operates a smelting facility at Tiwai Point near Invercargill. The smelter produces aluminium by the electrolytic reduction of alumina according to the reduction equation $2Al_2O_3 + 3C \rightarrow 4Al + 3CO_2$. This reaction which is called the Heroult-Hall process, is carried out as a continuous process in reduction cells constructed of an outer steel shell and a lining of refractory bricks. A carbon cathode is placed in the floor of the cell and carbon anodes are suspended above the cell on cast iron yokes. A very high direct current of approximately 190,000 amps is passed between the anode and cathode through a bath of molten cryolite at 960^0C which provides the electrical conductivity. The alumina is feed into the cryolite bath at regular intervals from a hopper that is located above the cell. The carbon required in the reduction reaction is provided by the carbon anode blocks which gradually reduce in size over a period of approximately twenty-seven days. When a block becomes too small, it is replaced by a new block. As the aluminium is produced it sinks to the bottom of the cell. Each day approximately 1260kgs of molten aluminium are tapped from the cell into a crucible by a vacuum siphoning system. A crucible is a large steel bucket lined with refractory bricks. Each crucible can tap the aluminium from three cells.

The cells are laid out in four lines. Three of the lines are each approximately 600 metres long and consist of 204 cells grouped into four tapping bays each made up of 51 cells. The fourth line 300 metres long was installed more recently and is made up of 48 cells of a newer technology in one tapping bay. All the cells in a tapping bay are tapped once each day and produce seventeen (or sixteen in the case of line 4) crucibles. Each bay is tapped either during the day shift or during the night shift. Once each crucible is filled with aluminium from three cells (always from the same tapping bay), it is transported from the reduction lines to furnaces in the Metal Products Division from where it is cast into finished products in the form of ingot or billet.

The purity of the aluminium varies from cell to cell depending on a number of factors including the age of the cell, the purity of the alumina feed and the manner in which the cell has been operated during its production life. The purity of the aluminium

declines gradually as contaminants in the form of iron, silicon, gallium and other chemicals increase until at some stage a decision is made to cease production in the cell. The cell is then taken off-line and rebuilt before being brought back into production some days later. Each cell is assayed regularly to determine the percentage of aluminium, iron, silicon, gallium and other chemicals. Because high purity aluminum commands a premium price on the metals market, it is important that aluminium tapped into a crucible from high purity cells is not contaminated by tapping from low purity cells into the same crucible.

This Case Study describes the development of a set partitioning optimisation model to batch or group triples of cells so that the total value of metal produced is maximised. The main aim is to minimise the dilution of high purity (high value) metal by low purity (low value) metal. The optimised batches tend to group high purity cells together and leave the lower purity cells to be batched with other lower purity cells. Numerical results show that significant improvements in excess of 15% can be achieved in the value of metal by carefully batching cells.

In the following section of this paper, we will describe an optimisation model for cell batching and discuss the formulation of a natural objective to measure the solution quality. We will then discuss aspects of the solution process and in particular outline how integer solutions can be derived from continuous LP relaxation solutions. Some numerical results will be then be presented to show the benefits that mathematics can bring.

An Optimisation Model for Cell Batching

The cell batching optimisation can be formulated naturally as a set partitioning problem (SPP) which can be written as

$$\text{minimise } z = \mathbf{c}^T\mathbf{x}, \quad A\mathbf{x} = \mathbf{e}, \quad x_j = 0 \text{ or } 1$$

where A is a 0-1 matrix and $\mathbf{e}^T = (1,1,...,1)$. Because cells in different tapping bays can never be tapped into the same crucible, the cell batching optimisation problem for each tapping bay can be considered independently of the other bays. For each tapping bay, the 51 constraints or rows of A correspond to cells in the tapping bay and ensure that each cell appears in exactly one batch or crucible. The columns of A represent all possible triples of cells (i.e. batches) which could be tapped into the same crucible. Each column then has exactly three nonzero unit values. For example, a batch made up of cells 1, 3 and 6 would be represented in the model by a column with zeros everywhere except for unit values in rows 1, 3 and 6. In general then the elements a_{ij} are defined as

$$\begin{aligned} a_{ij} &= 1 \text{ if cell } i \text{ is included in batch } j \text{ and} \\ &= 0 \text{ otherwise.} \end{aligned}$$

In this basic unrestricted form of the model there are ${}^{51}C_3$ or 20825 columns or variables which can be easily enumerated. A solution of this SPP will be made up of exactly seventeen variables at unit value (representing the chosen batches) and all other variables will have zero value.

Spread Limitations on Cell Batches

Because each tapping bay is approximately 300 metres long, it is not practical to tap cells that are far apart into the same crucible. The actual tapping process involves the

use of a gantry crane to carry the crucible and the human operator walks along the line from cell to cell. To avoid requiring the operator to walk large distances to fill each crucible, the usual practice is to tap cells that are within some specified maximum distance apart on the line. This is referred to as the spread of the batch. Spread can be defined simply as the difference between the maximum and minimum cell number in the batch. So three adjacent cells have a minimal spread of two. A batch made up of cells 1, 3 and 6 would have a spread of 5. In the enumeration of the columns or batches in the SPP model, the spread for each batch can be calculated. If the spread exceeds a specified limit, the batch can be rejected and not included in the model. Alternatively, the batch could be included in the model but marked as having a spread exceeding the maximum spread. During the optimisation process, these batches would be ignored unless the maximum spread was increased sufficiently. Such an increase could easily be included in a post-optimal investigation. In Section 3 we comment further on such investigations.

Batches exceeding the maximum spread can also be considered in a more useful manner. By adding an additional generalised SPP constraint to the basic SPP model, we could permit a limited number of batches with excessive spread to be included in the solution. This reflects the management view that a small number of batches (e.g. one or two of the seventeen batches) with excessive spread can be tapped provided they generate a sufficiently improved optimal solution when compared to the optimal solution using only batches that are within the maximum spread limit. Typical spread limits might be a maximum spread of 5 but up to two batches with a maximum spread of 10. In the enumeration of columns in the SPP model, all batches with a spread up to 10 would be generated but all those batches with spread between 6 and 10 would contribute to the additional constraint with a right-hand-side of 2. These spread restrictions which are included either implicitly (i.e. no spread exceeding 10) or explicitly (i.e. limited spread exceeding 6) significantly reduce the total number of variables in the SPP to something less than about 3000.

Alloy Codes and a Cell Batching Objective Function

A natural objective for the cell batching optimisation can be based on some estimate of the market value of the aluminium. This will obviously reflect the purity of each batch. Batch purity can be calculated as a simple weighted average of the known cell purities which make up the batch where the weights reflect the weight of metal tapped from each cell. Although in practice the actual cell tapping weights do vary a little, it is reasonable to assume before the tapping takes place that the tapping weights will be constant at 1260kgs for lines 1, 2 and 3 and 1480kgs for line 4. Given constant tapping weights and the cell assay values for aluminium, iron, silicon and gallium, the batch values are calculated as simple averages.

Table 1 defines eighteen aluminium alloy codes with their corresponding minimum aluminium percentage and maximum percentages for silicon, iron and gallium and an estimate of the corresponding alloy premium value. Generally speaking as the aluminium percentage increases (with corresponding decreases in the silicon, iron and gallium percentages) the alloy premium value increases rapidly. While this Table refers more particularly to finished product values cast in the particular alloy codes, we can classify the batch purity using the alloy specifications and then use the premium value as an objective coefficient for the batch. The negative premium values for the first three alloy codes reflect the fact that these grades of aluminium usually require

purification by mixing with a higher grade metal and thus result in an actual loss of premium.

Table 1
Alloy Code Specifications and Premiums

Code	Min Al%	Max Si%	Max Fe%	Max Ga%	Premium
AA????	0.000	1.000	1.000	1.000	-50.00
AA150	99.500	0.100	0.300	0.100	-40.00
AA160	99.600	0.100	0.300	0.100	-25.00
AA1709	99.700	0.100	0.200	0.100	0.00
AA601E	99.700	0.100	0.080	0.100	40.00
AA601G	99.700	0.100	0.080	0.100	40.00
AA185G	99.850	0.054	0.094	0.014	15.00
AA190A	99.900	0.054	0.074	0.014	45.00
AA190B	99.900	0.050	0.050	0.014	50.00
AA190C	99.900	0.035	0.037	0.012	110.00
AA190K	99.900	0.045	0.055	0.034	100.00
AA191P	99.910	0.030	0.045	0.010	120.00
AA191B	99.910	0.030	0.027	0.012	139.00
AA192A	99.920	0.030	0.040	0.012	140.00
AA194A	99.940	0.020	0.040	0.007	150.00
AA194B	99.940	0.034	0.034	0.010	180.00
AA194C	99.940	0.022	0.027	0.009	200.00
AA196A	99.960	0.020	0.015	0.010	260.00

Off-line Cells

When cells are taken off-line to be rebuilt, they will not be tapped. This implies that one or more of the seventeen batches from the tapping bay will include fewer than three cells. It is important that the optimised solution determine which batches should be composed of fewer cells. One approach would be to generate all possible batches involving one and two cells and include them in the SPP model when the tapping bay has off-line cells. This results in a very large increase in the number of variables and causes further computational problems during the solution process. A much more attractive approach is to simply treat off-line cells as having a zero tapping weight. The off-line cells are then permitted to appear anywhere in the cell order during the enumeration of batches. In other words, the off-line cell can appear in any triple of cells without affecting either the spread calculation or the batch chemical composition. For example, if cell 50 is off-line, then a batch made up of cells 1, 2 and 50 would have a spread of 1 and the chemical composition would be determined entirely by cells 1 and 2. If this batch were included in the optimal solution, it would be interpreted as a batch involving just cells 1 and 2. However the SPP constraint for cell 50 would have been satisfied by this variable. The advantage of this approach is that all batches remain triples of cells including all triples involving off-line cells and the SPP model (at least for lines 1 to 3) is always made up with 51 cell constraints.

The Solution Process

The SPP model is solved by first solving the LP relaxation problem in which the integer restrictions are relaxed. The LP solution is trivial. With little extra effort it is possible to report on a sequence of LP solutions which gradually include variables with wider and wider spreads. During this initial optimisation phase, batches with spreads

exceeding the maximum permitted spread are ignored or equivalently, the right-hand-side for the additional constraint described in Section 2.1 is set to zero thus preventing batches with excessive spread from contributing to the solution.

Handling Excessive Spread Batches

It is also simple to quantify the benefits of permitting a small number of batches with spreads exceeding the specified maximum spread as discussed in Section 2.1. After completing the initial LP solution, the right-hand-side for the additional constraint can be increased slowly to identify the potential benefits of using a limited number of batches with wider spread. All of these calculations are performed using the LP relaxation. While the LP solutions do exhibit some evidence of natural integer structure, most solutions involve fractional variables that must be forced to integer values using a branch and bound algorithm.

Constraint Branching in the Cell Batching Optimisation Model

The natural constraint branch (see Ryan and Foster, 1981) for this SPP is defined by any pair of cells that are not always included together in batches in the fractional solution. It is easy to show using balanced matrix theory (see Ryan and Falkner, 1988) that such a pair of cells must always exist in any fractional solution. The binary constraint branch can then be imposed by requiring on the one-side of the branch that the two cells always appear together in a batch and on the zero-side of the branch that the two cells always appear in different batches. The implementation of the branch involves removing sets of variables from the descendant LPs. On the one-side, batches in which the cells do not appear together are removed (effectively by setting their corresponding variable upper bounds to zero) and on the zero-side batches in which the cells appear together are removed.

While this constraint branching strategy is particularly effective, it is true that after imposing a sequence of constraint branches, the LP can become infeasible because a cell becomes isolated from its neighbours in such a way that no feasible set of batches can include the cell. When this happens it usually results in a long sequence of fathoming infeasible nodes and the branch and bound process can take a long time to find a feasible integer solution. For this reason we have implemented a heuristic integer allocation process at each node in the branch and bound tree in order to find integer solutions more quickly.

Integer Allocation Heuristics

Because of the special structure of the underlying SPP model in this application, it is easy to create infeasible LPs at nodes in the branch and bound tree. To avoid the computational problems that this causes, we have implemented an integer allocation process which is applied at each fractional node in the branch and bound tree including the root node. The process is heuristic in that it attempts to force an integer solution from the fractional solution by making a sequence of greedy decisions. First, all variables (batches) at value one in the fractional solution are fixed at that value. We then find the first cell which is not yet included in a batch and search amongst all variables (including those which are nonbasic) for the least cost batch including that cell and two other cells which are also yet to be covered. If no such variable can be found, the search is abandoned and the integer allocation fails. The allocation process can be applied using any order of the cells. In our implementation we apply the allocation process considering the cells in increasing and also in decreasing order.

Table 2

Optimised and Default Solutions for Tapping Bay 2AW

Optimised solution for tapping bay 2AW: [Cells 352 to 402] [Dataset d271197]												
Cell 15(-366) is off-line												
Cell 41(-392) is off-line												
A total of 51 cells found for bay 2AW; 2 cells off-line												
Generating with MINSPREAD 4; MAXSPREAD 10; MAXSPREADRHS 2												
Model size: 52 constraints; 2237 variables												
Bay	#	Weight	%AL	%SI	%FE	%GA	Code	Spread	Prem	Cells		
2AW	1	3840	99.800	0.044	0.095	0.016	AA1709	2	0	352	353	354
2AW	2	3840	99.810	0.040	0.105	0.016	AA1709	2	0	355	356	357
2AW	3	3840	99.737	0.041	0.167	0.015	AA1709	4	0	358	360	362
2AW	4	3840	99.847	0.042	0.080	0.014	AA601E	4	40	359	361	363
2AW	5	3840	99.760	0.041	0.132	0.016	AA1709	3	0	364	365	367
2AW	6	2560	99.835	0.034	0.079	0.015	AA601E	3	40	-366	399	402
2AW	7	3840	99.813	0.036	0.101	0.016	AA1709	2	0	368	369	370
2AW	8	3840	99.823	0.038	0.096	0.015	AA1709	3	0	371	372	374
2AW	9	2560	99.840	0.034	0.080	0.015	AA601E	2	40	373	375	-392
2AW	10	3840	99.820	0.038	0.095	0.016	AA1709	2	0	376	377	378
2AW	11	3840	99.833	0.035	0.087	0.016	AA1709	2	0	379	380	381
2AW	12	3840	99.837	0.037	0.079	0.015	AA601E	5	40	382	384	387
2AW	13	3840	99.817	0.037	0.093	0.016	AA1709	3	0	383	385	386
2AW	14	3840	99.837	0.033	0.080	0.016	AA601E	7	40	388	389	395
2AW	15	3840	99.753	0.075	0.126	0.015	AA1709	3	0	390	391	393
2AW	16	3840	99.843	0.037	0.080	0.015	AA601E	3	40	394	396	397
2AW	17	3840	99.827	0.036	0.092	0.016	AA1709	3	0	398	400	401
Integer objective: 240 (Branch and Bound time of 12.82 seconds)												
Default 2 spread solution for tapping bay 2AW: [Cells 352 to 402]												
2AW	1	3840	99.800	0.044	0.095	0.016	AA1709	2	0	352	353	354
2AW	2	3840	99.810	0.040	0.105	0.016	AA1709	2	0	355	356	357
2AW	3	3840	99.770	0.041	0.153	0.013	AA1709	2	0	358	359	360
2AW	4	3840	99.813	0.041	0.094	0.016	AA1709	2	0	361	362	363
2AW	5	2560	99.760	0.040	0.121	0.016	AA1709	1	0	364	365	-366
2AW	6	3840	99.793	0.038	0.119	0.016	AA1709	2	0	367	368	369
2AW	7	3840	99.827	0.035	0.093	0.015	AA1709	2	0	370	371	372
2AW	8	3840	99.830	0.038	0.090	0.015	AA1709	2	0	373	374	375
2AW	9	3840	99.820	0.038	0.095	0.016	AA1709	2	0	376	377	378
2AW	10	3840	99.833	0.035	0.087	0.016	AA1709	2	0	379	380	381
2AW	11	3840	99.820	0.040	0.090	0.016	AA1709	2	0	382	383	384
2AW	12	3840	99.833	0.034	0.083	0.016	AA1709	2	0	385	386	387
2AW	13	3840	99.820	0.035	0.087	0.016	AA1709	2	0	388	389	390
2AW	14	2560	99.715	0.093	0.144	0.015	AA1709	2	0	391	-392	393
2AW	15	3840	99.847	0.037	0.077	0.015	AA601E	2	40	394	395	396
2AW	16	3840	99.837	0.035	0.083	0.015	AA1709	2	0	397	398	399
2AW	17	3840	99.837	0.035	0.088	0.016	AA1709	2	0	400	401	402
Default 2-spread objective: 40												

The process is remarkably effective in that it often produces integer solutions with objective values very close to the LP upper bound value. Such an integer solution often becomes the best bound that fathoms the remaining live nodes in the branch and bound tree and the solution process terminates.

Some Numerical Results

We report here (see Table 2) some results that illustrate the performance of the cell batching optimisation and compare the results with a so-called default order solution.

Table 3
Solution Summary by Alloy Codes for all Tapping Bays

		Optimised solution		Default solution	
Code	Premium	# crucibles	premium	# crucibles	premium
AA160	-25	3.00	-75.00	3.00	-75.00
AA1709	0	27.67	0.00	73.67	0.00
AA190K	100	1.00	100.00	5.33	533.33
AA192A	140	31.67	4433.33	40.00	5600.00
AA194B	180	22.00	3960.00	16.00	2880.00
AA194C	200	1.67	333.33	1.00	200.00
AA196A	260	9.00	2340.00	3.00	780.00
Totals		219.00	16011.67	219.00	12998.33
Percentage increase in total default premiums 23.18%					

The default order solutions, based on batching cells in the natural sequence such as (1,2,3), (4,5,6), (7,8,9) etc, actually form the basis of the manual solution method in which obvious poor value batches are modified locally by changing the allocations of cells to nearby batches to improve the solution. This is a difficult process for human decision making and even experienced operators are unlikely to produce optimal decisions. Table 2 includes cost values based on the alloy code premium for each batch. It can be seen that in tapping bay 2AW, the optimised solution produces 6 batches with premium value of 40.0 (alloy code AA601E) while the default solution produces just one batch with this premium value. Notice also that batches 12 and 14 in the optimised solution involve spreads of 5 and 7 respectively. All other batches have spreads not exceeding 4. The two off-line cells (366 and 392) appear in batches 6 and 9 respectively as negative cell numbers. The weights, spreads and chemical compositions of these two batches ignore the off-line cells.

In Table 3 we give a comparison of the overall results of the optimised cell batching applied to all tapping bays. The optimised results can be compared with the corresponding overall results of the default solution. The results are reported in terms of the number of batches (i.e. crucibles) produced in each alloy code and the total premiums generated by those batches. The optimised solutions show an improvement of approximately 23% over the default solutions.

Conclusions

While the cell batching optimisation has produced significant improvements in the value of metal produced from the reduction lines, this problem really forms part of a

larger production scheduling problem at the smelter. The full problem involves first a decision about which products from the order book to produce during the day. This production must then be scheduled on the furnaces and casting machines in the Metal Products Division. A furnace production schedule is made up of periods during which the furnace is filled with suitable batches from the reduction lines followed by periods during which the chemical composition of the furnace metal is adjusted and stabilised before the required product is cast. Given a production schedule for each furnace, the allocation of batches from the cell batching optimisation to furnace fills must be decided. This particular problem is not trivial since it implies a time sequencing for the actual production of the batches. There are particularly important constraints on the production of batches that result from the limited capacities of the gantry cranes used during the tapping process. In the cell batching optimisation, these constraints were ignored as was the time sequencing of production of the batches. We are currently investigating this further scheduling problem to determine a feasible production sequence for the optimal batches to match a given furnace schedule.

Case 2: Optimised Crew Scheduling for Air New Zealand Ltd

In the mid 1980s, the scientific literature contained relatively few papers documenting the successful application of optimisation methods in the solution of airline crew scheduling problems although some heuristic methods were being applied. In fact, stories circulating in the airline industry at that time suggested that a number of larger airlines had tried and failed to implement optimisation based crewing systems. In fact, when I first approached Air New Zealand in the early 1980s to ask if I could obtain information about their crew scheduling problems, I was informed by a senior manager that he was aware of the failures of other airlines and he asked me what made me think I could solve crewing problems for Air New Zealand. I remember responding that I didn't know if I could solve the problems but I simply wanted to find out about the problems and obtain some data to try solving them. In 1984, Air New Zealand agreed to provide information about their planning or Tour of Duty (Pairings) problem for us to use in an Honours project for a student in Engineering Science at the University of Auckland. This initial project, involving a small part of the domestic (or internal) problem, was undertaken by Michelle Kunath. The results we presented to Air New Zealand at the completion of the six-month project provided the basis of a most productive and successful collaboration between the University of Auckland and Air New Zealand. Over the intervening period of more that 16 years, optimisation methods have been developed and implemented to solve all aspects of Air New Zealand's crew planning and rostering problems.

In 1984, all crewing decisions were made manually and the airline used no Operations Research (OR) techniques. Today the airline is totally dependent on state of the art optimisation based computer systems in the areas of crew planning and rostering. The airline now employs eight staff with backgrounds in OR. In this Case Study, we document the transition from dependence on manual methods to dependence on mathematical optimisation methods in New Zealand's national airline.

The Crew Scheduling Problems

Airline crew scheduling involves two distinct processes of Planning and Rostering. The Planning process (also referred to as the Pairings problem) involves the construction of a minimum cost set of generic Tours of Duty (ToDs) or pairings which cover all relevant flights (sectors) in an airline flight schedule. Each Tour of Duty

begins and ends at a crew base and consists of an alternating sequence of duty periods and rest periods with duty periods including one or more sectors. At Air New Zealand, domestic (or National) ToDs (on B737 aircraft) are between one and three days in duration while for International airline operations (on B767 and B747 aircraft), ToDs can be up to fourteen days in duration. New ToDs are constructed for each new flight schedule and for day to day variations in an underlying schedule. It should be noted that the construction of ToDs does not involve any consideration of the crewmembers who will actually perform the ToDs.

The Rostering process involves the allocation of ToDs to each crew member of a rank so that all flights are crewed with the correctly qualified crew complements and each crew member has a legal feasible line of work over a given roster period. At Air New Zealand, domestic rosters are built over a fourteen-day roster period while international rosters cover a twenty-eight day period. Roster construction must take into account activities rostered in the previous roster period, which carry over into the current period. From a crew point of view, it is also important to provide high quality rosters that satisfy crew requests and preferences as much as possible.

During the past two or three decades airlines have invested heavily in the development of techniques to solve their crew scheduling problems. The main reasons for this focus on crew scheduling can be identified in the following major factors.

- Reducing Aircrew costs
  Aircrew costs are one of the largest operating costs faced by an airline (second only to fuel). However, the crewing problems of Planning and Rostering are very large and hard to solve. Manual and heuristic-based solution methods will almost never find minimum cost solutions due to the very large number of alternative solutions and the complex nature of the crew scheduling rules. Because small improvements in solution quality return large dollar savings, the use of optimisation techniques to solve the Planning problem has been the primary focus of much research within the Airline Industry for many years.

- Reducing solution time
  The time taken to create solutions manually meant that only one solution could be developed, and alternative proposals could not be evaluated quickly or accurately. A shorter planning cycle means that the airline can respond quickly to changes in the market and capitalise on opportunities. The flight schedule can be changed much more frequently and solutions must be continually updated. The challenge is to maintain crew productivity and the quality of the crewing solution while the flight schedule is changed from day to day.

- Compliance
  Crew scheduling is constrained by many complex and conflicting rules, further exacerbated by time-zone changes, daylight savings, and foreign currencies. Solutions must comply with legislative, contractural and operational rules. Airlines are required to have systems in place to ensure that rules are not violated. The complexity of the rules is also a factor in the time taken to solve these problems and manually produced solutions are not able to guarantee compliance.

- Reducing costs to construct and maintain the crew scheduling process.
  The complex nature of the rules and the experience required to achieve consistently high-quality manual solutions mean that staff changes could result in increased crewing costs.

Over the past sixteen years considerable research and development of underlying optimisation methods for crew scheduling has been undertaken at the University of Auckland in collaboration with Air New Zealand. This research has resulted in the development of seven optimisation-based computer systems to solve all aspects of both the Planning and the Rostering processes for both the National and the International Airlines. It should be noted that each business problem is characterised by quite unique aspects that prevent the development of a single common optimisation solver. The following Table shows the implementation dates of each system. The term "Technical Crew" refers to pilots (Captains and First Officers).

	Flight Attendants		Technical Crew	
	Planning	Rostering	Planning	Rostering
National	1986 revised 1997	1993	1986 revised 1997	1998
International	1998	1989 revised 1996	1996	1994

These systems incorporate state-of-the-art mathematical optimisation technology and provide Air New Zealand with sophisticated crewing solvers. In 1989 when the International flight attendant rostering system was implemented, Air New Zealand knew of no other airline worldwide with an implemented rostering systems based on optimisation and even today, few airlines use optimisation based rostering techniques. The optimisation solvers are fully integrated into other information systems at Air New Zealand. Further details of each of these systems and their integration are given below.

Many scheduling problems can be formulated mathematically using a Set Partitioning model. From a technical point of view, this specially structured zero-one integer linear programme which has relatively few constraints but a very large number of variables model is a. Both the Planning and the Rostering problems of airline crew scheduling can be formulated as set partitioning problems with special structure. Research conducted at the University of Auckland, in collaboration with Air New Zealand, has resulted in major breakthroughs in the solution of very large instances of set partitioning models which occur in practical applications. By recognising the special model structure and incorporating it in the solution methods, it is possible to solve optimisation problems that just ten years ago were considered far too difficult to solve. In particular, practical instances of the Rostering problem can be very large but it is still possible to produce high quality solutions. These important technical developments involving

- limited subsequence matrix generation
- constraint branching strategies for integer programming
- resource constrained shortest path column generation and
- anti-degeneracy and steepest edge pricing strategies

have been described by Ryan and Falkner (1988), Ryan and Osborne (1988), Ryan (1992).

The Tours of Duty Planning Model

In the basic ToD planning model, each column or variable in SPP corresponds to one possible ToD that could be flown by some crew member. Each constraint in SPP corresponds to a particular flight and ensures that the flight is included in exactly one ToD. The elements of the A matrix can then be defined as

$a_{ij} = 1$ if the jth ToD (variable) includes the ith flight (constraint)

$= 0$ otherwise.

The value of c_j, the cost of variable j, reflects the dollar cost of operating the jth ToD. The calculation of c_j values is specified by the particular problem being considered but usually includes the cost of paid hours (both productive and unproductive), ground transport, meals and accommodation, and the cost of passengering crew within the ToD. Many authors (see Rubin, 1973 and Wedelin, 1995) model the ToD planning problem using the set covering formulation in which the equality constraints are replaced by "greater than or equals" constraints. The overcover of a flight permitted by the set covering constraint can be interpreted as passengering of the excess crew cover. This formulation results in fewer variables but it has the major disadvantage of making it difficult to model accurately the rules and costs associated with passengering crew. For example, passengering duty time limits are generally longer than operating limits. This cannot be correctly modelled using set covering constraints. In the Air New Zealand applications, set partitioning constraints are used which allows passengering to be accurately modelled. This results in additional columns which explicitly include passengering flights with $a_{ij} = 0$. It is also important to note that each column or ToD must correspond to a feasible and legal sequence of flights which satisfies the rules specified in civil aviation regulations or employment contracts or agreements. These rules or constraints can be thought of as being implicitly rather than explicitly satisfied in the ToD planning model. For example, rules imposing limits on total work time and rest requirements are embedded in the variable generation process.

The ToD planning model is usually augmented with additional constraints that permit restrictions to be imposed on the number of ToDs included from each crew base. Because these constraints typically have non-unit right-hand-side values, we describe the ToD planning model as a generalised set partitioning model.

In Air New Zealand, the ToD planning model is formed and solved independently for each crew type and the flights they operate. We will identify further specific variations and extensions of the basic ToD planning model in the detailed discussion of each of the Air New Zealand systems.

The Rostering Model

The rostering problem involves the construction of a LoW for each crew member in a rank so that each ToD is covered by the correct number of crew members from that rank. For each crew member we can generate a set of many LoWs from which exactly one must be chosen.

The rostering problem can also be modelled mathematically using a generalised version of the set partitioning model. Assuming there are p crew members and t ToDs, the model is naturally partitioned into a set of p crew constraints, one for each crew member in the rank, and a set of t ToD constraints corresponding to each ToD which

must be covered. The variables of the problem can also be partitioned to correspond to the feasible LoWs for each individual crew member. The A matrix of the rostering set partitioning model is a 0-1 matrix partitioned as

$$A = \begin{bmatrix} C_1 & C_2 & C_3 & \cdots & C_p \\ L_1 & L_2 & L_3 & \cdots & L_p \end{bmatrix}$$

and $C_i = \mathbf{e}_i\mathbf{e}^T$ is a (p x n_i) matrix with $\mathbf{e}_i$ the i^{th} unit vector and $\mathbf{e}^T = (1,1,...,1)$. The n_i LoWs for crew member i form the columns of the (t x n_i) matrix L_i with elements l_{jk} defined as $l_{jk} = 1$ if the k^{th} LoW for crew member i covers the j^{th} ToD and $l_{jk} = 0$ otherwise. The A matrix has total dimensions of m x $\sum_{i=1}^{p} n_i$ where $m = p + t$. The right-hand-side vector $\mathbf{b}$ is given by $b_i = 1$, $i = 1,...,p$ and $b_{p+i} = r_i$, $i = 1,..., t$ where r_i is the number of crew members required to cover the i^{th} ToD. We refer to the first p constraints as the "crew constraints", and the next t constraints as the "ToD constraints".

The cost vector $\mathbf{c}$ is chosen to reflect the relative "cost" of each LoW. Since most airlines do not use optimisation systems for rostering, there is no obvious or traditional measure which can be used to discriminate among feasible solutions in an optimisation. We define particular rostering objectives in the discussion of each of the specific rostering systems developed at Air New Zealand. Typically the rostering objective reflects either the "preferential bidding by seniority" (PBS) or the "equitable" rostering philosophy.

The rostering model has a special structure which deviates from pure set partitioning in that the right-hand-side vector is not unit valued and some constraints need not be equalities. The crew constraints of the A matrix also exhibit a generalised upper bounded structure which is not commonly found in set partitioning.

Crewing Systems at Air New Zealand

National (Domestic) Planning

The original Planning system for the National Airline covering both Flight Attendants and Technical crew was developed in 1984 and 1985 and implemented as a mainframe computer system in 1986. The system remained in production essentially in its original form until 1997 when it was replaced by improved optimisation methodology implemented on a Unix workstation. The current system generates optimised ToDs for all crew ranks and for three crew bases in Auckland, Wellington and Christchurch. It is also able to produce "fully-dated" solutions.

National (Domestic) Rostering

While involving relatively small crew ranks (at least compared to the International Airline), these problems are probably the most difficult of all the Air New Zealand optimisation problems to solve because of their combinatorial complexity. Two previous attempts to solve the problems in the late 1980s were unsuccessful but the problem for Flight Attendants was finally solved in 1993 by Dr Paul Day in his PhD research sponsored by Air New Zealand (see Day, 1996 and Day and Ryan, 1997). In 1998, the same solution methodology was adapted to produce Technical crew rosters under quite different operating rules. These two unique systems are now fully

integrated into the Air New Zealand Genesis Rostering System and produce rosters of excellent quality for all crew ranks and all crew bases in less than four person days. Previous manual rostering methods involved 6 roster builders and took two weeks to complete the roster build. The actual optimisation runs themselves take less than one hour in total.

International Flight Attendant Rostering

This problem, involving 1500 flight attendants in four crew ranks, is the largest problem solved at Air New Zealand. The original system was implemented in 1989 and was revised to incorporate column generation methods in 1996. At the time of its implementation in 1989, the optimised solution demonstrated that it was possible to construct rosters with a 5% reduction in the number of flight attendants and at the same time, significantly improve the quality of the rosters from a crew point of view. The development and implementation involved representatives of the Flight Attendant Union who defined the issues of roster quality which are incorporated in the optimisation. The current system also incorporates a language assignment optimisation step (Waite, 1995) which ensures that flight attendants with relevant language qualifications are assigned ToDs requiring those language skills. This aspect of Flight Attendant rostering has important commercial benefits to Air New Zealand in that many of its passengers, particularly from Asia and Europe, are non-English speaking.

International Technical Crew Rostering

International technical crews in most airlines world-wide are rostered by systems based on preferential bidding by seniority (PBS). The algorithms are generally based on greedy sequential heuristic roster construction methods. PBS involves crew members bidding for work or days off and rosters are then constructed by satisfying as many bids as possible but considering crew members strictly in seniority order within the crew rank. During 1992 and 1993, a new optimisation model and solution method for PBS was developed (see Thornley, 1993). The solution method incorporates a unique "squeeze procedure" which violates the bids of more junior crew members in order to satisfy the bids of more senior crew members. This guarantees that the maximum number of bids can be satisfied in seniority order. Heuristic methods used by other airlines are unable to provide such a guarantee. The PBS system was implemented in 1994 and is now fully integrated into the Genesis Rostering System at Air New Zealand.

International Technical Crew Planning

Following the completion of his Masters research on the topic, Andrew Goldie (see Goldie, 1995) implemented the Technical Crew Planning system for Air New Zealand International in 1996. The system automatically generates "third pilot" ToDs which allow duty periods to be extended by including a third pilot on some relevant sectors. This feature is believed to be unique since we understand that Planning systems used by other airlines construct such ToDs in a subsequent step.

International Flight Attendant Planning

The International Flight Attendant Planning problem is a particularly difficult problem in that flight attendants are qualified to operate on all aircraft types. The added complexity arises because each aircraft type requires different numbers of crew. For example, a full B747 crew may split after a B747 sector and part of the crew may fly a B767 sector in their next duty period. The remaining part of the B747 crew could fly as

passengers or could be combined with other crew members to make up a crew for some other sector. This crew splitting complication does not occur for Technical crew who are qualified to fly just one aircraft type. An optimisation solver for International Flight Attendant Planning has been developed by Chris Wallace in his PhD research. This system is again unique in that it automatically permits crew splitting. No other known Planning system incorporates this feature.

Implementation and Integration Issues

On-site development by a small team of developers, working closely with the users, has been central to the successful implementation of these systems. The complex rules-bound nature of the industry requires detailed understanding, and the optimisation solvers must be developed with constant reference to the planners and rosterers.

The optimisation solvers are able to find many solutions of similar dollar value, some of which are preferred by the users, and much effort has been spent developing control mechanisms for the users to interact with the solutions and so "shape" the solutions produced. For example, users may wish to fix a particular subset of a solution, or prevent particular undesirable characteristics from being included in a solution. Similar mechanisms have been developed to handle changes to inputs to the problem. For example, if a flight is re-timed in the flight schedule, the optimiser will minimise the changes required from the previous solution to the new solution.

The ToD optimisers are integrated into a purpose-built PC-based user interface, also developed as part of the project. The system receives flight schedule data for both proposed and published flight schedules in industry-standard formats from Airflight, the Schedules Management tool supplied by The Sabre Group. The solutions produced by the ToD optimisers can be viewed graphically, using a tool specifically developed for the purpose. Solutions may be electronically uploaded into the Air New Zealand Genesis Rostering System.

The Genesis Rostering System, which has been developed independently of this project, is used by all Rostering Staff to manage the construction of rosters for aircrew. It has replaced existing mainframe and PC-based systems and manual methods, previously used for roster construction. It provides a common user interface for managing ToDs, crew pre-assignments, training and crew requests. Genesis passes data to the Rostering optimisers where the roster is constructed, before it uploads and displays the optimised rosters in the graphical interface.

There are two important aspects associated with the implementation of sophisticated crewing systems: the first is concerned with the effects of new technology and the second is concerned with issues of technology transfer.

The introduction of high-technology mathematical optimisers has changed the nature of the jobs and the key competencies of staff in the Planning and Rostering areas. Staff now manage data and processes associated with the construction of ToDs and rosters rather than simply constructing ToDs or rosters. The optimisers allow staff to concentrate on meeting specific business requirements which might include provision of training or leave, or meeting special requirements of management or crew, or incorporating last-minute changes.

Because of the sophisticated nature of the mathematical optimisation models and methods, it is important that management, the crew and the planning and rostering staff

develop trust and confidence in the solution methods and the quality of the solutions. This can only be achieved through close collaboration with these affected groups. The development of trust and confidence has been a major objective of this project since it is an essential requirement of successful technology transfer from a research and development environment to a production environment. We believe that this objective has been fully achieved in this project.

Economic and Other Benefits

The crew scheduling optimisers provide real dollar benefits to Air New Zealand, by directly reducing the cost of crewing in areas such as the total number of crew required and the number of hotel bed-nights, meals, and other expenses which must be paid to crew away overseas. Each optimisation application has reduced the costs of constructing and maintaining the crewing solution for the flight schedule. Over the past ten years, Air New Zealand's aircraft fleet and route structure has increased significantly in size yet the number of staff required to solve the crew scheduling problem has reduced significantly from a total of 27 in 1987 to 15 today.

A conservative estimate of the dollar savings (in New Zealand dollars) from the crew scheduling optimisers has been calculated as NZ$15,655,000 per annum. It is interesting to contrast the estimated total savings of NZ$15.6 million per annum with the estimated total development costs over 15 years of approximately NZ$2 million. It is also interesting to contrast the estimated total savings of NZ$15.6 million per annum with the 1999 net operating profit for the Air New Zealand group of NZ$133.2 million (excluding one-off adjustments).

In addition to the direct dollar savings, many intangible benefits are also provided by these optimisation systems.

1. Using the crewing optimisers, high quality solutions can be produced in a matter of minutes, compared with two or more days to create a manual solution. For example, a B767 pilot ToD problem can be optimally solved in approximately 60 minutes, while the B747-400 pilot ToD problem can be solved in less than 5 minutes on a Unix workstation. The international flight attendant rostering problem involving 550 crew in one rank is solved in less than six hours.

2. Crew schedulers can now focus on data preparation and validation, and the interpretation and evaluation of solutions. Their role has changed from the mechanical process of ToD and roster construction to one of an analyst.

3. There is a reduced dependence on highly skilled staff with an intimate knowledge of the employment contracts and scheduling rules. These contracts and rules are now embedded in the crewing optimisers.

4. The crewing optimisers can be used to investigate strategic decisions for crewing such as the evaluation of proposed rule changes and their cost impacts, and basing studies to determine ideal crew numbers at crew bases.

5. The relatively short build times for ToDs and rosters enables solutions to be produced much closer to the day of operation and with reduced lead times. This enables the company to accommodate late schedule changes and reduces rework.

6. The ToD optimisers make it possible to provide accurate and reliable feedback to the schedules planning group about proposed schedules and the impact on profit

from a crewing point of view. This information was very difficult to provide manually because of the time taken to produce manual ToD solutions.

7. Before the implementation of the national ToD optimiser in 1986, Air New Zealand operated a fixed six monthly winter and summer schedule. Now Air New Zealand is able to operate a much more flexible schedule that varies from week to week and allows the company to respond quickly to market opportunities.

8. The ToD optimisers also make it possible to repair solutions quickly when small changes are made to the schedule. The cost to maintain solutions is kept to a minimum.

9. The rostering optimisers reflect crew defined roster quality measures and crew are encouraged to identify "soft rules" to further improve roster quality. The involvement of the crew in the development of systems that affect their lifestyle is a very important benefit that cannot be provided adequately by manual systems.

10. The rostering optimisers have delivered significantly improved levels of crew request and bid satisfaction. Over 80 percent of all legal international flight attendant requests for ToDs and days off are consistently achieved, with even higher achievement of requests for the national pilot and flight attendant rostering systems.

11. The rostering optimisers can accurately identify roster infeasibility and can minimise the level of infeasibility. They can also identify any days on which there are insufficient crew.

12. Improved passenger service has been achieved directly though the use of the international flight attendant languages optimisation. Air New Zealand's high levels of customer service have been recognised through recent awards including the "Globe Award for the Best Airline to the Pacific" awarded by top British travel industry newspaper Travel Weekly in January 2000. Air New Zealand has received this award in four of the last five years. Air New Zealand was also ranked first for inflight service by the AB Road Airline Survey, in August 1999.

13. The crewing optimisers provide a guarantee that important legislative and contractual rules are satisfied. In many of these situations, Air New Zealand is required to demonstrate compliance through an audit procedure. Manual systems are unable to guarantee compliance without time consuming checking.

References

Day, P. R., & Ryan, D. M. (1997). Flight attendant rostering for short–haul airline operations, *Operations Research*, 45(5), 649–661.

Day, P. R. (1996). *Flight Attendant Rostering for Short–haul Airline Operations*, PhD Thesis, University of Auckland.

Goldie, A. P. (1995). *Optimal Airline Crew Scheduling Using Dynamic Column Generation*, Masters Thesis, University of Auckland.

Ryan, D. M. (1992). The solution of massive generalised set partitioning problems in aircrew rostering, *J. Op. Res. Soc.* 43, 459–467.

Ryan, D. M., & Falkner, J. C. (1988). On the integer properties of scheduling set partitioning models, *European Journal of Operations Research*, 35, 442–456.

Ryan, D. M., & Foster, B. A. (1981). An integer programming approach to scheduling, in A. Wren (Ed) *Computer Scheduling of Public Transport*, North Holland, 269–280.

BEST COPY AVAILABLE

Ryan, D. M. & Osborne, M. R. (1988). On the solution of highly degenerate linear programmes, *Math Prog.* 41, 385–392.

Thornley, M. E. (1993). *Crew Rostering Under A Seniority Preferential Bidding Environment Using Column Generation*, Masters Thesis, University of Auckland.

Tuck, S. (1997). *Optimal Cell Batching for New Zealand Aluminium Smelters Tiwai Point Facility*, Year Four Project, Dept of Engineering Science, University of Auckland.

Waite, J. (1995). *Language Assignment For Flight Attendant Rostering*, Masters Dissertation, University of Auckland.

U.S. Department of Education
Office of Educational Research and Improvement (OERI)
National Library of Education (NLE)
Educational Resources Information Center (ERIC)

NOTICE

Reproduction Basis

[X] This document is covered by a signed "Reproduction Release (Blanket)" form (on file within the ERIC system), encompassing all or classes of documents from its source organization and, therefore, does not require a "Specific Document" Release form.

[ ] This document is Federally-funded, or carries its own permission to reproduce, or is otherwise in the public domain and, therefore, may be reproduced by ERIC without a signed Reproduction Release form (either "Specific Document" or "Blanket").